arXiv:1611.04199v1 [cs.LG] 13 Nov 2016

Realistic risk-mitigating recommendations via inverse classification

Michael T. Lash and W. Nick Street

University of Iowa, Iowa City, USA
{michael-lash, nick-street}@uiowa.edu

Abstract. Inverse classification, the process of making meaningful perturbations to a test point such that it is more likely to have a desired classification, has previously been addressed using data from a single static point in time. Such an approach yields inflated probability estimates, stemming from an implicitly made assumption that recommendations are implemented instantaneously. We propose using longitudinal data to alleviate such issues in two ways. First, we use past outcome probabilities as features in the present. Use of such past probabilities ties historical behavior to the present, allowing for more information to be taken into account when making initial probability estimates and subsequently performing inverse classification. Secondly, following inverse classification application, optimized instances' unchangeable features (e.g., age) are updated using values from the next longitudinal time period. Optimized test instance probabilities are then reassessed. Updating the unchangeable features in this manner reflects the notion that improvements in outcome likelihood, which result from following the inverse classification recommendations, do not materialize instantaneously. As our experiments demonstrate, more realistic estimates of probability can be obtained by factoring in such considerations.

1 Introduction

Individual-centric data mining and machine learning paradigms have the potential to enrich the lives of millions of people by first making accurate predictions as to their future outcomes and then secondly recommending implementable courses of action that make a desired future outcome more likely. Consider semi-annual medical physicals that assess the current status of an individual's well-being, for example, Patient 8584, taken from our experiments. We have information on three semi-annual medical examinations for this patient and know that, sometime between the second and third visit, this person was diagnosed with cardiovascular disease (CVD). After learning a predictive model and evaluating the patient at the first visit, their predicted probability of CVD was 22%. Their predicted probability at the second visit was 36%. Merely observing this individual's progression towards this negative life-altering event is informative, but unhelpful: crucial steps detailing how such risk can be reduced are needed.

In this work we propose using inverse classification, the process of making recommendations using a machine learning method to optimally minimize

or maximize the probability of some outcome, in conjunction with longitudinal data. Specifically, we want to minimize patient 8584's probability of CVD beginning from the first medical exam we have on record. This is the first contribution of this work: to methodologically incorporate longitudinal data into the inverse classification process.

Considerations we make, leading to more realistic assessments of risk and more meaningful recommendations, constitute our second and third contributions. Specifically, as demonstrated by Patient 8584, past medical visits have a bearing on feature patient visits and, ultimately, on whether or not there is an unfavorable outcome; cumulative actions made over a period of time led to this person experiencing an adverse event. Therefore, we propose to incorporate the predicted risk at previous visits as a feature in future visits. Doing so makes estimates of current risk more realistic and allows the inverse classification process to make recommendations that account for past behaviors.

Furthermore, when making recommendations using inverse classification, and subsequently estimating the probability of a particular outcome, we can further obtain a more realistic estimate of risk by using what are referred to as the immutable feature-values observed at the proceeding medical visit. Immutable features are attributes that can't be changed, such as one's age. We include these values from future visits in the estimation of risk to make the assessment more realistic. The inclusion of such future feature values reflects the fact that changes made by an individual, such as Patient 8584, are not made instantaneously (contribution three), but take time to implement.

The rest of the paper proceeds by first discussing related work in Section 2, our proposition on the incorporation of longitudinal data into the inverse classification process, and how we, specifically, make our three contributions, in Section 3. In Section 4 we discuss our experiment process, parameters, and results, demonstrating these three contributions (and a fourth tertiary contribution, discussed later). Section 5 concludes the paper.

2 Related Work

Emergent data mining and machine learning research involving **longitudinal data** is focused on methodologically leveraging such data, as well as the specific domains in which such methods can be employed. In [14] the authors explore deep neural networks that, with minimal preprocessing, learn a mapping from patients' lab tests to over 130 diseases (multi-task learning). In [15] unsupervised learning methods are applied to longitudinal health data to learn a disease progression model. The model can subsequently be used to aid patients in making long-term treatment decisions. These works exemplify the way in which models can be learned to aid in predicting disease [14] and in forecasting disease progression [15]. In this work we examine how (1) coupling longitudinal data with predictive models can make disease risk estimation more accurate and (2) how predictive models that incorporate historical risk and past behavior can be used

to make recommendations that optimally minimize the likelihood of developing a certain disease.

Inverse classification methods are varied in their approach to finding optimal recommendations, either adopting a greedy [1, 5, 9, 17] or non-greedy formulation [3, 7, 8, 12]. Past works also vary in their implementation of constraints that lead to more realistic recommendations, either being completely unconstrained [1, 5, 17], or constrained [3, 7, 8, 9, 12]. In this work, we adopt the formulation and framework related by [7, 8] which accounts for (a) the features that can and cannot be changed (e.g. age cannot be changed, but exercise levels can), (b) varying degrees of change difficulty (feature-specific costs) and (c) a restriction on the cumulative amount of change (budget). As in [7], we implement a method that avoids making greedy recommendations while still accounting for (a), (b) and (c).

3 Inverse Classification

In this section we begin by discussion some preliminary notation, followed by our inverse classification framework, which we subsequently augment to account for past risk and missing features.

3.1 Preliminaries

Let $\mathbf{X}_v = [\mathbf{x}_{v_1}, \ldots, \mathbf{x}_{v_n}]^\top \in \mathbb{R}^{n \times p}$ denote a matrix of training instances and $\mathbf{y}_{v+1} = [y_{v_i}, \ldots, y_{v_n}]^\top \in \{0,1\}^n$ their corresponding labels at visit v, where $v = 1, \ldots, V$. Here, individual visits v can be viewed as discrete, no-overlapping time units, where $\mathbf{X}_v$ is observed at a discrete time unit and $\mathbf{X}_{v+1}$ is observed at the next discrete time unit $v+1$. We note here that X_v is observed at one visit and the event of interest $\mathbf{y}_{v+1}$ is observed at the next discrete time unit, namely $v+1$. We do this to reinforce the fact that we are interested in modeling how the current state of instances $\mathbf{X}_v$ results in some not-to-distant future outcome. Additionally, let $\mathbf{X}_{v+1} \subseteq \mathbf{X}_v$ hold. That is, the instances at $v+1$ or also present at v, reflecting that these datasets are longitudinal.

Let $f_v : \mathbb{R}^p \to \mathbb{R}$ denote a trained classifier that has learned the mapping $\mathbf{X}_v \to \mathbf{y}_{v+1}$. Once trained, $f_v(\mathbf{x})$ can be used to make a prediction as to the short-term outcome of test instance $\mathbf{x}$ at visit $v+1$. Here, f_v can be any number of classification functions.

With these preliminaries in mind we ultimately want to leverage longitudinal data to accomplish two things: **(1)** to obtain realistic recommendations for $\mathbf{x}_v$ and **(2)** to use estimated risks from previous visits as predictive features in the present to further improve risk estimates of $\mathbf{x}_v$.

3.2 Inverse classification

To address **(1)**, an optimization problem must be formulated. Consider

$$\begin{aligned} &\min_{\mathbf{x}} f_v(\mathbf{x}_v) \\ &\text{s.t. } \mathbf{c}_v^\top \mathbb{I}(\mathbf{x}_v - \bar{\mathbf{x}}_v) \leq B_v \\ &\quad l_j \leq x_j \leq u_j, \text{ for } j = 1, \ldots, p \end{aligned} \tag{1}$$

where $\mathbf{x}$ is being optimized, $\bar{\mathbf{x}}$ is the original instance, $\mathbf{c}$ is a cost vector, $\mathbb{I}$ is a signature matrix where $\mathbb{I}_{jj} \in \{-1, 1\}$, B is a specified budget and l_j, u_j are imposed lower and upper bounds, respectively. To control the direction of feature optimization we set

$$l_i/u_i = \begin{cases} \mathbf{x}_{D_j} & \text{if } \mathbb{I}_{jj} = 1/\mathbb{I}_{jj} = -1 \\ \min(0, \mathbf{x}_{D_j})/\max(1, \mathbf{x}_{D_j}) & \text{otherwise} \end{cases} \tag{2}$$

which has the effect of increasing (decreasing) features when $\mathbb{I}_{jj} = 1$ (-1).

In a realistic setting, such as the medical domain of our experiments below, not all $\mathbf{x}_j$ are mutable. To reflect this, we partition the features into to two non-overlapping sets, U and C, representing the immutable and mutable features, respectively. Not all mutable features can be changed directly, however. Some such features are altered indirectly. To reflect this, we further partition C into D and I, reflecting those that are directly mutable and those that are indirectly mutable, respectively. Additionally, we define a function $\phi : \mathbb{R}^{|U|+|D|} \rightarrow \mathbb{R}^{|I|}$ that estimates the indirectly mutable features values using the features U and D. Incorporating these distinctions into (2) gives us

$$\begin{aligned} &\min_{\mathbf{x}_C} f_v(\mathbf{x}_{v_U}, \phi(\mathbf{x}_{v_U}, \mathbf{x}_{v_D}), \mathbf{x}_{v_D}) \\ &\text{s.t. } \mathbf{c}_v^\top \mathbb{I}(\mathbf{x}_v - \bar{\mathbf{x}}_v) \leq B_v \\ &\quad l_{v_j} \leq x_{v_j} \leq u_{v_j}, \text{ for } j \in D \end{aligned} \tag{3}$$

where we are optimizing only the changeable features D.

Past Risk To address **(2)**, we propose using $r_{v-1} = f_{v-1}(\mathbf{x}_{v-1})$ and then incorporating r_{v-1} as a feature of $\mathbf{x}_v$ such that $\mathbf{x}_v = [(r_1, \cdots, r_{v-1}, \mathbf{x}_{v_U}), \mathbf{x}_{v_I}, \mathbf{x}_{v_D}]^\top$. We believe that the inclusion of past risk as features in the present will lead to more accurate probability estimates and, as a result, better inform the inverse classification process. Our proposition to include past risk as predictive features in the present can be observed graphically in Figure 1, albeit in simplified form.

Missing Feature Estimators Practically speaking, features S at one v are not always present in a subsequent time period. Namely, $S_{v=k} \neq S_{v=q}$ where $k, q \in \{1, \ldots, V\}$. To overcome this issue, we propose defining the full set of features at $v = 1$, such that $S_{v=q} \subseteq S_{v=1}$ and $S_{v=q} \cap S_{v=1} = S_{v=q}$. During

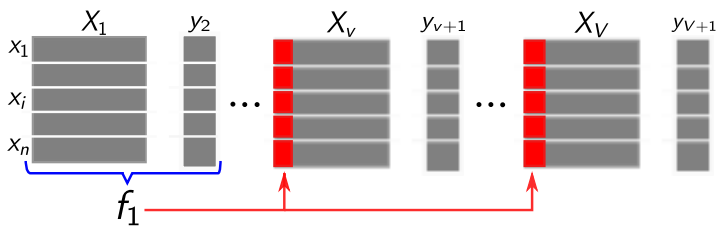

Fig. 1: Using predicted risk at a previous visit as a feature in future visits.

subsequent visits q, we propose to use an estimator, using the subset of features available at q, but learned using $v = 1$ data, to impute these values. To reflect this let $\psi_m : \mathbb{R}^{p_{v=q}} \rightarrow \mathbb{R}$, for $m \in M_{v=q}$, where ψ_m is either a regressor or a classifier, depending on the feature being estimated. This results in $S_{v=1} \equiv \{S_{v=q}, S_{M_{v=q}}\}$

Taking past risk and the estimation of missing features into account, the optimization problem can be reformulated to

$$\begin{aligned} \min_{[\mathbf{x}_C, \mathbf{x}_{C_M}]} \; & f_v([\mathbf{x}_U, \mathbf{x}_{U_M}], \phi([\mathbf{x}_U, \mathbf{x}_{U_M}], [\mathbf{x}_D, \mathbf{x}_{D_M}]), [\mathbf{x}_D, \mathbf{x}_{D_M}]) \\ \text{s.t. } & c_v^\top \mathbb{I}([\mathbf{x}, \mathbf{x}_M] - [\bar{\mathbf{x}}, \bar{\mathbf{x}}_M]) \leq B_v \\ & l_j \leq x_j \leq u_j, \text{ for } j \in D \end{aligned} \tag{4}$$

where we abuse the notation by letting U_M and D_M denote the missing features for each of the respective U, D feature sets and assume that each of these missing feature sets proceeds those that are known. Also note that past risk features have been merged with $\mathbf{x}_{v_U}$.

Implementing Recommendations and Risk Estimates Performing Equation (4) produces $\mathbf{x}_v^*$, an optimized version of $\mathbf{x}_v$ whose risk has been minimized. In the real world, however, it takes time to implement the feature updates that actually take $\mathbf{x}_v$ to $\mathbf{x}_v^*$. To reflect this, we replace $\mathbf{x}_{v_U}^*$ with $\mathbf{x}_{v+1_U}$ and update $\mathbf{x}_{v_I}^* = \phi(\mathbf{x}_{v+1_U}, \mathbf{x}_{v_D}^*)$. In other words, we use the immutable features at the next discrete time unit, and a new estimate of the indirectly mutable, based on these updated immutable features, along with the optimized mutable features, to obtain a more realistic $\mathbf{x}_v^*$. We fully outline the inverse classification process in Figure 2.

Finally, performing inverse classification in the manner we have outlined requires an optimization mechanism, of which there are many. To further zero in on a suitable method, we outline an assumption and the desirable qualities we'd like the optimization method to have. First, as in [7], we assume that f_v is differentiable and it's gradient is L-Lipschitz continuous. That is, $\|\nabla f_v(\mathbf{x}) - \nabla f_v(\mathbf{x}')\| \leq L\|\mathbf{x} - \mathbf{x}'\| \forall \mathbf{x}, \mathbf{x}' \in \mathbb{R}^{|D|}$. If we assume that f_v is linear, and observing that our constraints in Equation (4) are also linear, the optimization can be performed efficiently. However, we do not wish to make this prohibitive

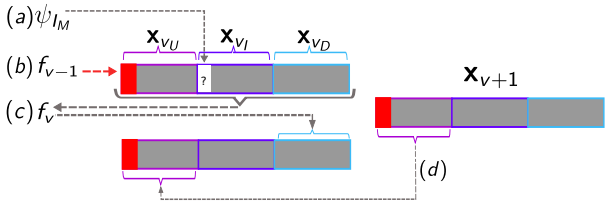

Fig. 2: The inverse classification process using longitudinal instances. (a) Missing features are estimated (in this example there is a missing indirectly mutable feature). (b) Past risk is used as an immutable feature in the present. (c) Inverse classification is performed. (d) The immutable features at $v+1$ are used to estimate the new level of risk for the optimized test instance.

assumption, although we do wish to solve the problem efficiently. Therefore, we elect to use the projected gradient method [6, 11], which has been shown to have a convergence rate of $O(\frac{1}{t})$ for nonlinear problems whose constraints are linear.

4 Risk Mitigating Recommendations

In this section we outline the data used in our experiments, our experiment process, and results.

4.1 Data

Our data is derived from the Atherosclerosis Risk in Communities (ARIC) study [2]. The investigation began in 1987 wherein 16000 individuals from four different communities were initially examined by physicians. Individual assessments entailed the thorough documentation of a variety of patient-characterizing attributes. These components can be categorized as demographic (e.g., age), routine (e.g., weight), lab-based (e.g., blood glucose levels), and lifestyle-based (e.g., amount of exercise). Subsequent follow-up exams were conducted on a bi-annual basis thereafter.

ARIC data is freely available, but does require explicit permission-to-use prior to acquisition. Additionally, some processing and cleansing must be done. For the sake of reproducibility, the code that takes these data from their acquisition state to the final state used in our experiments is provided at github.com/michael-lash/LongARIC

.

We use three sets of data in our experiments, representing three bi-annual physician visits. Our target variable is defined based on outcomes observed at $v+1$. These outcomes pertain to cardiovascular disease (CVD) events. We define a CVD event to be one of the following: probable myocardial infarction (MI), definite MI, suspect MI, definite fatal coronary heart disease (CHD), possible fatal CHD, or stroke. Patients having such an event recorded at their $v+1$

examination have $y_{v+1} = 1$, whereas patients not having any such events at $v + 1$ have $y_{v+1} = 0$. Patients for whom one of these events are observed at a previous visit are excluded from subsequent visit datasets. We do this to ensure that we are continuing, at each visit, to learn a representation that is consistent with two-year risk of CVD. Therefore, $\mathbf{X}_3 \subseteq \mathbf{X}_2 \subseteq \mathbf{X}_1$. Table 1 summarizes the number of patients, features/missing features, and positive instances at each of the visits.

Dataset	Instances	Feats/Missing	$y_{v+1} = 1$
X_1	12223	122/0	232
X_2	11057	98/24	249
X_3	9883	74/48	231

Table 1: Dataset descriptors.

4.2 Experiment Information

Evaluative Process The process of conducting our experiments is carefully crafted such that no data used in making recommendations is ultimately used in evaluating their success (final probability estimates). As such, each of the X_v datasets are partitioned into two parts at random. One part is used to train an f_v that is used for the inverse classification process. The second partition is further partitioned into 10 parts. One part is used as a test set for which recommendations are obtained, while the other nine are used to train a validation model that evaluates the probability of CVD resulting from the recommendations obtained. By maintaining partition separation, and iteratively cycling between the role each partition plays, recommendations can be obtained for all X_v instances. The process is more definitively outlined in [7].

Experiment Parameters In our experiments we choose to use RBF SVMs [4] as our f_v, which can be trained by solving the dual optimization problem

$$\max_{\alpha \in \mathbb{R}^n} \sum_{i=1}^{n} \alpha_i - \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_i \alpha_j y^i y^j k(\mathbf{x}^i, \mathbf{x}^j) \tag{5}$$
$$\text{s.t.} \sum_{i=1}^{n} \alpha_i y^i = 0 \text{ and } 0 \leq \alpha_i \leq C \text{ for } i = 1, 2 \ldots, n$$

where $k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x}-\mathbf{x}'\|^2}{2\sigma^2}\right)$ is the Gaussian kernel, $\sigma > 0$ and $\|\cdot\|$ represents the Euclidean norm in $\mathbb{R}^p$. Practically speaking, any number of other kernels could have been selected. We elect to use these as they were observed to have good inverse classification performance in [7]. Furthermore, by employing

Platt Scaling [13] we can directly learn a probability space that is more easily interpreted.

Secondly, we elect to use kernel regression [10, 16] as our indirect feature estimator ϕ_v, which is given by

$$\mathbf{x}_{v_I} = \frac{\sum_{i=1}^{n} k([\mathbf{x}^i_{v_D}, \mathbf{x}^i_{v_U}], [\mathbf{x}_{v_D}, \mathbf{x}_{v_U}])\mathbf{x}^i_{v_I}}{\sum_{i=1}^{n} k([\mathbf{x}^i_{v_D}, \mathbf{x}^i_{v_U}], [\mathbf{x}_{v_D}, \mathbf{x}_{v_U}])}, \tag{6}$$

where the $k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x}-\mathbf{x}'\|^2}{2\sigma^2}\right)$ Gaussian kernel and the value $\sigma > 0$ is selected based on cross-validation. We elect to use this method as the estimation is based on point similarity, which is consistent with our assumption that similar points will have similar CVD probabilities.

Experiments Using this evaluative procedure and outlined learning methods, we propose three experiments. Two of these experiments directly address how the inclusion of past risk and $\mathbf{x}_{v+1_U}$ affect probability estimation, while the third is somewhat tertiary to this objective.

Experiment 1: We demonstrate that more realistic estimates of probability, following the application of the inverse classification process, can be obtained by leveraging longitudinal data. We do so by first performing inverse classification on each of the three X_v and estimate the resulting probabilities independent of one another. Then, we do the same procedure, except we use $\mathbf{x}_{v+1_U}$ instead of $\mathbf{x}_{v_U}$, and $\mathbf{x}_{v_I} = \phi(\mathbf{x}_{v+1_U}, \mathbf{x}^*_{v_C})$ instead of $\mathbf{x}_{v_I} = \phi(\mathbf{x}_{v_U}, \mathbf{x}^*_{v_C})$ to estimate the probabilities. This is related by (d) in Figure 1.

Experiment 2: We examine the collective impact of including past risk as a predictive feature in the present, related by (b) in Figure 1, in conjunction with $\mathbf{x}_{v+1_U}$ in performing inverse classification and assessing the resulting outcome probabilities.

Experiment 2 is outlined by Figure 3.

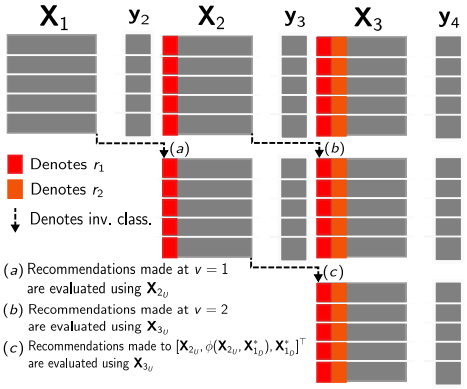

Fig. 3: The inverse classification process for longitudinal instances.

Experiment 3: We would like to see whether the use of a learning method is better suited to estimating missing features in subsequent visits, or whether a simple carry-forward procedure is better. In other words, can we simply use the known feature values at a previous visit $v-1$ to estimate missing features at v, or is it better to use a learning method? To test what method works best we randomly selected two continuous features and one binary feature that are known in $v = 2$. We then trained various classifiers/regressors (where appropriate), using the known $v = 2$ features, but with $v = 1$ data, along with the carry-forward procedure, to make estimates. We then observed what performed the best in terms of MSE (continuous) or AUC (binary).

4.3 Results

In this section we present the results of our three experiments, albeit somewhat out of order. We begin by showing the results of Experiment 3, wherein we discovered the best method of estimating missing features at future visits. We present this first, as the results are used in practice in the subsequent two experiments.

Method \ Feat/type	Alcohol/cont	Statin Use/bin	Hematocrit/cont
Carry	50.55	.579	7.55
RBF SVM	47.65	.50	9.31
Lin SVM	47.65	.50	9.31
CART	30.97	.569	2.37
kNN	**29.37**	.50	12.81
Log Reg	NA	**.984**	NA
Ridge	**29.42***	NA	**1.24**

Table 2: Missing feature estimation (Experiment 3) results given in MSE for continuous features and AUC for binary.

Table 2 shows the results of applying the described carry-forward procedure and a host of learning algorithms to the three randomly selected features. In all cases at least one learning method outperforms the carry-forward procedure. For continuous features, we observe that Ridge regression works well, and that kNN is either the best or the worst model. Therefore, we estimate continuous features using Ridge regression. The binary feature is best estimated by logistic regression. We use these selected learning methods to estimate missing features at $v = 2, 3$.

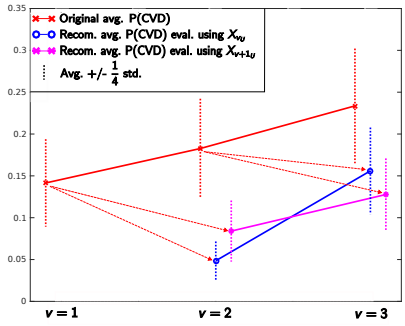

Fig. 4: Results from Experiment 1.

Figure 4 shows the results of **Experiment 1**. Here we observe the average predicted probability and $\frac{1}{4}$ of one standard deviation across the three visits, shown in red, and the predicted probability after applying the inverse classification process at $v = 1, 2$: blue shows evaluation of the probability using $\mathbf{x}_{v_U}$ and magenta shows evaluation of the probability using $\mathbf{x}_{v+1_U}$. We first observe that the inverse classification process was successful in reducing P(CVD) at both visits ($v = 1, 2$). Secondly, we observe that probabilistic results differ between blue and magenta, which directly supports our hypothesis that accounting for recommendation implementation time leads to different probability estimates.

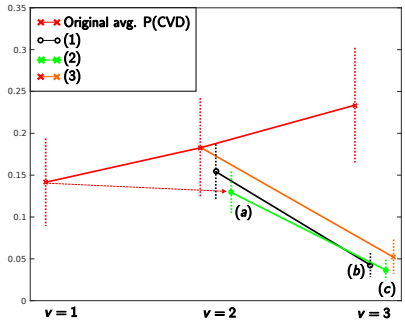

Fig. 5: Results from Experiment 2.

Figure 5 shows the results of **Experiment 2**. Here, we present results in a similar manner to Experiment 1, but with some slight differences; the means and $\frac{1}{4}$ of one standard deviation are represented as they previously were, and red still denotes the original predicted probability. **(1)**, represented in black, starts with the predicted probability at $v = 2$, taking into account $r_v = 1$. The inverse

classification process is then applied, leading to the predicted probability at $v = 3$, represented by (b) – this corresponds to (b) in Figure 3. **(2)**, represented in green, begins with optimized instances from $v = 1$, corresponding to (a) in Figure 3, and ending with (c) in the same figure. **(3)**, shown in orange, is the predicted probability at $v = 3$ taking both r_1 and r_2 into account.

The first red, black, and orange data points in Figure 5 show the average predicted probability of CVD, taking past risk into account (i.e., they parallel what is shown in red, differing only in that past risk has been incorporated). Comparing the black point at $v = 2$ and orange point at $v = 3$ we see that there are marked differences between the unaccounted for risk predictions (red) and those that take past risk into account: past risk incorporation has lead to lower probability estimates. The lower estimates are intuitive, as our learning method has become more certain of those individuals who will not develop CVD – low past risk likely indicates no immediate threat of developing a disease that takes years to manifest. As Table 1 shows, $\approx 2\%$ of instances develop CVD at each v, which supports the lower average probability estimates.

The same holds true for the two green and one black data points in Figure 5. These points represent averages (and $\frac{1}{4}$ of one standard deviation $\pm$ the mean) obtained after applying inverse classification, accounting for past risk and using $\mathbf{x}_{v+1_U}$ in the estimates. We observe only slight improvements in mean CVD probability after performing the process, versus the more extreme values observed in Figure 4. Additionally, we can see that (b) and (c) in Figure 5 are very similar – (c) is .4% lower on average. This is likely the result of diminishing returns. Once appropriate lifestyle adjustments have been made, making further adjustments are unlikely to be as beneficial.

5 Conclusions

In this work we proposed two ways in which longitudinal data can be fused with an existing inverse classification framework to arrive at more realistic assessments of risk. First, we used past risk as an immutable predictive feature in the present. Second, after making recommendations, to reflect the fact that implementation and benefit are not observed instantaneously, we used the unchangeable features from the next discrete time period to estimate probability improvements. In our experiments we noticed that the inclusion of such factors resulted in different probability estimates than those obtained without. Future work in this area may benefit from the use of methods that are capable of taking all prior and current features into account, as well as future unchangeable features, which may cumulative help in making the inverse classification process as precise as possible.

References

1. Aggarwal, C.C., Chen, C., Han, J.: The inverse classification problem. Journal of Computer Science and Technology 25(May), 458–468 (2010)

2. ARIC Investigators and others: The Atherosclerosis Risk in Communitities (ARIC) study: Design and objectives (1989)
3. Barbella, D., Benzaid, S., Christensen, J., Jackson, B., Qin, X.V., Musicant, D.: Understanding support vector machine classifications via a recommender system-like approach. In: Proceedings of the International Conference on Data Mining. pp. 305–11 (2009)
4. Boser, B.E., Guyon, I.M., Vapnik, V.N.: A training algorithm for optimal margin classifiers. In: Proceedings of the Fifth Annual Workshop on Computational Learning Theory. pp. 144–152. ACM (1992)
5. Chi, C.L., Street, W.N., Robinson, J.G., Crawford, M.A.: Individualized patient-centered lifestyle recommendations: An expert system for communicating patient specific cardiovascular risk information and prioritizing lifestyle options. Journal of Biomedical Informatics 45(6), 1164–1174 (2012), http://dx.doi.org/10.1016/j.jbi.2012.07.011
6. Ghadimi, S., Lan, G.: Stochastic first- and zeroth-order methods for nonconvex stochastic programming. SIAM Journal on Optimization 23, 2341–2368 (2013)
7. Lash, M.T., Lin, Q., Street, W.N., Robinson, J.G.: A budget-constrained inverse classification framework for smooth classifiers. arXiv preprint; arxiv:1605.09068 (2016), https://arxiv.org/abs/1605.09068
8. Lash, M.T., Lin, Q., Street, W.N., Robinson, J.G., Ohlmann, J.: Generalized inverse classification. arXiv preprint; arxiv:1610.01675 (2016), https://arxiv.org/abs/1610.01675
9. Mannino, M.V., Koushik, M.V.: The cost minimizing inverse classification problem : A genetic algorithm approach. Decision Support Systems 29(3), 283–300 (2000)
10. Nadaraya, E.a.: On estimating regression. Theory of Probability & Its Applications 9(1), 141–142 (1964)
11. Nesterov, Y.: Gradient methods for minimizing composite objective function. Mathematical Programming, Series B 140, 125–161 (2013)
12. Pendharkar, P.C.: A potential use of data envelopment analysis for the inverse classification problem. Omega 30(3), 243–248 (2002)
13. Platt, J., et al.: Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. Advances in Large Margin Classifiers 10(3), 61–74 (1999)
14. Razavian, N., Marcus, J., Sontag, D.: Multi-task prediction of disease onsets from longitudinal lab tests. arXiv preprint arXiv:1608.00647 (2016)
15. Wang, X., Sontag, D., Wang, F.: Unsupervised learning of disease progression models. In: Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 85–94. ACM (2014)
16. Watson, G.S.: Smooth regression analysis. The Indian Journal of Statistics, Series A 26(4), 359–372 (1964)
17. Yang, C., Street, W.N., Robinson, J.G.: 10-year CVD risk prediction and minimization via inverse classification. In: Proceedings of the 2nd ACM SIGHIT symposium on International health informatics - IHI '12. pp. 603–610 (2012), http://dl.acm.org/citation.cfm?id=2110363.2110430